arXiv:1412.3898v1 [cs.IR] 12 Dec 2014

ILCR: Item-based Latent Factors for Sparse Collaborative Retrieval

Lu Yu
Alibaba Research Center for Complexity Sciences
Hangzhou Normal University
Hangzhou, China
coolluyu@gmail.com

Junming Huang
Web Sciences Center
University of Electronic Science and Technology of China
mail@junminghuang.com

Chuang Liu
Alibaba Research Center for Complexity Sciences
Hangzhou Normal University
Hangzhou, China
liuchuang@hznu.edu.cn

Zi-Ke Zhang
Alibaba Research Center for Complexity Sciences
Hangzhou Normal University
Hangzhou, China
zhangzike@gmail.com

ABSTRACT

Interactions between search and recommendation have recently attracted significant attention, and several studies have shown that many potential applications involve with a joint problem of producing recommendations to users with respect to a given query, termed *Collaborative Retrieval* (CR). Successful algorithms designed for CR should be potentially flexible at dealing with the sparsity challenges since the setup of collaborative retrieval associates with a given $query \times user \times item$ tensor instead of traditional $user \times item$ matrix. Recently, several works are proposed to study CR task from users' perspective. In this paper, we aim to sufficiently explore the sophisticated relationship of each $query \times user \times item$ triple from items' perspective. By integrating item-based collaborative information for this joint task, we present an alternative factorized model that could better evaluate the ranks of those items with sparse information for the given query-user pair. In addition, we suggest to employ a recently proposed scalable ranking learning algorithm, namely BPR, to optimize the state-of-the-art approach, *Latent Collaborative Retrieval* model, instead of the original learning algorithm. The experimental results on two real-world datasets, (i.e. *Last.fm*, *Yelp*), demonstrate the efficiency and effectiveness of our proposed approach.

Categories and Subject Descriptors

H.3.3 [**Information storage and retrieval**]: Information Search and Retrieval—*Information filtering*

General Terms

Algorithms, Experimentation

Keywords

collaborative filtering, information retrieval, matrix factorization, recommender systems

1. INTRODUCTION

The task of *Collaborative Retrieval* (or CR for short), producing recommendations to a particular individual with respect to a certain query describing her instant interests, arises from the rapidly emerging needs to explore relative information from huge resource. For example, a user may hope web application to retrieve some songs or movies similar to his/her recent favourites depending on a given query. Collaborative retrieval roots in personalized recommender systems and information retrieval systems but is largely different from a simple combination of those two isolated tasks. Expanding a user-item preference matrix in a recommender system or a query-document relevance matrix in an information retrieval system to a query-user-item affinity tensor, a collaborative retrieval system suffers more severely from data sparsity, especially in the scenario that the explosive growth of information in recent decades has built a huge candidate space.

To tackle the sparsity problem in collaborative retrieval, an effective and efficient method is required to fulfill the $query \times user \times item$ tensor. Tensor factorization models [23, 29] work well to extend collaborative filtering to recommend items in a tensor setup, but usually designed for special recommendation tasks like tag, mobile app, etc. According to [32], most of them can not fully fit the collaborative retrieval task. Additional source is another considerable option to reduce the pain of sparsity, including integrating side information like social networks [20] or trust-based network [10]. Nevertheless, the extra noise brought by those approaches severely affect sparse objects. In the other school, less noisy object relations are leveraged to supplement missing features of sparse objects, such as leveraging similar users [9] and similar items [27, 7]. An early exploration to incorporate object relation in collaborative retrieval is the *Latent Collaborative Retrieval* (LCR) model proposed by Weston *et al.* [32]. It leverages user-user similarity in modelling a $query \times user \times item$ tensor, in hope of reducing the pain of

sparsity. LCR partially addressed the sparsity problem but ignores the item side, which should not be negligible for two reasons. (1) Items suffer from the sparsity problem more severely than users, since items are usually observed with fewer features to support a feature-based or content-based algorithm. (2) Users are dynamic but items are relatively static, which makes user-user similarity less stable and reliable than item-item similarity, as reported in [19].

Inspired by hose intuitions, we propose in this paper an *Item-based Local Collaborative Retrieval* (ILCR) model, which leverage item-item similarity to solve the sparsity problem in collaborative retrieval. Specifically, we use a triple (q, u, a) to describe a fact that a user u prefers to an item a with content features correlated to query q. Viewing items as media vertices, we utilize item's neighbours with similar tastes to overcome the sparsity problem caused by the lack of content features. The ILCR model is trained with a recent proposed pairwise learning algorithm *Bayesian Personalized Ranking* (BPR) [22], which has been applied in many published recommendation tasks including tag recommendation [24], relation extraction [25], focused matrix factorization for advertisement [12]. As pointed out in [9, 31] and verified on a real dataset *Last.fm*[1] [5], BPR significantly reduces the runtime required to train a latent model like LCR or ILCR while keeping the same performance, compared with the training strategy *Weighted Approximate-Rank Pairwise* (WARP) that was previously used in LCR.

In summary, the main contributions of this paper include:

- Comparing with WARP learning algorithm, the application of BPR algorithm on optimizing the parameters of LCR can sharply increase the training speed and simultaneously preserve almost the same performance on evaluation metric.
- We consider the item-based collaborative information to sufficiently alleviate the sparsity problem caused by items' lack of content features, and subsequently propose the ILCR method.
- The experimental results on the *Last.fm* and *Yelp* datasets show that the proposed algorithm ILCR is superior to LCR model, especially, when dealing with a dataset containing a massive amount of items with sparse information.

2. COLLABORATIVE RETRIEVAL

The objective of CR can be simply defined as generating a personalized ranking list of items to fit a particular user's tastes with respect to a given query. To achieve this goal, the proposed approaches should clearly define a scoring function f to represent the relevance of a given triple (query,user,item) $\in \mathcal{Q} \times \mathcal{U} \times \mathcal{A}$, where $\mathcal{Q}$, $\mathcal{U}$, $\mathcal{A}$ denote the set of queries, users, items, respectively. In practice, only the top-k retrieved items could draw users' attention. Thereby, the learned scoring function f should promote users' interesting items to high position as much as possible for a particular query. Generally, the parameters of f can be derived from the training samples by optimizing a pre-defined ranking loss function.

[1]http://www.last.fm

We next briefly review the aspects of LCR model for CR task. Then, we detail the principles of the proposed model, namely ILCR.

2.1 Latent Collaborative Retrieval

The central idea of latent collaborative retrieval (LCR) is to represent each entity (e.g. user, query, item) as a n-dimensional feature vector. Analogous to MF-based approaches [16], the relationship between each pair entities can be measured by calculating the dot-product of them. Formally, LCR's parameter space includes matrices $S \in \mathbb{R}^{|\mathcal{Q}| \times n}$, $V \in \mathbb{R}^{|\mathcal{U}| \times n}$, and $T \in \mathbb{R}^{|\mathcal{A}| \times n}$, which denotes the feature matrix of queries, users, and items, respectively. To preciously evaluate a user's preference on a item w.r.t. a given query, LCR additionally allocates each user u an encoder matrix $U_u \in \mathbb{R}^{n \times n}$. The scoring function f of LCR model can be then given as follows:

$$f(q, u, a) = S_q U_u T_a^\top + V_u T_a^\top, \tag{1}$$

where S_q represents the row of S corresponding to query q, V_u is the row of V corresponding to user u, and T_a denotes the row of T corresponding to item a. Intuitively, the first term in Equation (1) distinguishes LCR from the proposed tensor models for context-aware CF[2] [1, 13, 24, 3] tasks. The second term independent of the query denotes users' basic preferences on items.

3. ITEM-BASED LATENT COLLABORATIVE RETRIEVAL

In this section, we will firstly describe our motivation behind the idea of Item-based LCR. Then we will show how to measure the relevance of $(query, user, item)$ triple from item-based perspective after graphically revisiting LCR method.

3.1 Motivation

In real applications, we often suffer from the problem of $asymmetric\ information$, that is to say, characteristics of a majority of items can not be fully expressed due to the lack of specific description. By contrast, massive amount of popular items are represented as adequate descriptive terms. In this circumstance, the retrieval systems would unfairly evaluate the ranks of those items with sparse information since current description can not cover the characteristics of them. In order to cope with this problem, many web applications, such as *delicious*[3], douban[4], integrate the function of *collaborative tagging* [11], which offers privileges for users to add free-style keywords to the shared content, e.g. books, movies, videos. To a certain extent, collaborative tagging enriches the content of items with descriptive terms for future filtering or search. However, there are still a huge amount of items with sparse information because of the lack of motivated users to share their tags. Therefore, the dataset that we have deeply involved mainly contains items either represented as adequate keywords, or oppositely with seldom keyword.

[2]The goal of context-aware CF is to produce personalized recommendations to a particular user by taking into account the contextual information, such as time, location, and so on. Therefore, the query could also be regarded as a context factor.

[3]https://delicious.com/

[4]http://www.douban.com

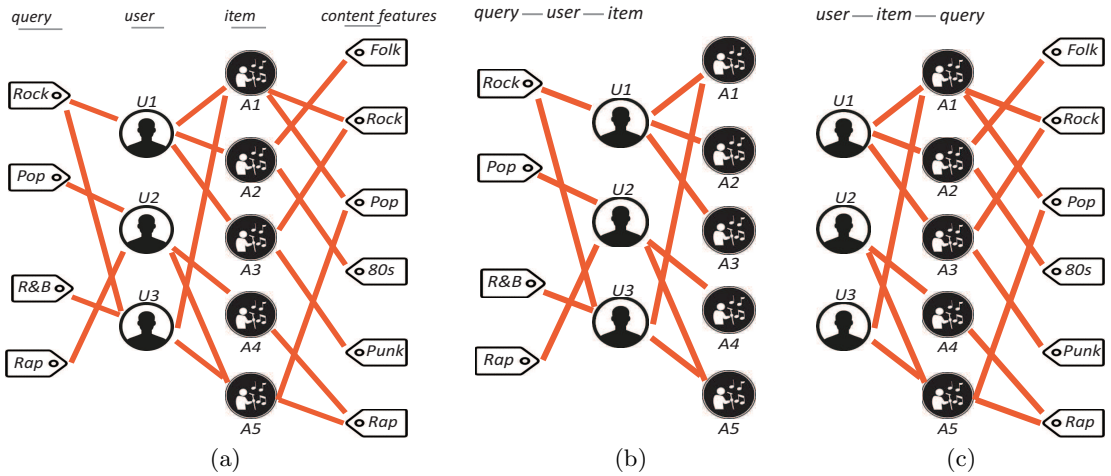

Figure 1: Graphical representation of the $user \times query \times item$ relationship.

To promote the performance of retrieval systems, we assume that the characteristics of those items with sparse information, denoted as *sparse items*, can be derived from their neighbours with abundant descriptive terms. Typically, one can firstly employ efficient tag recommendation approaches [8, 11] to pre-target sparse items with most possible keywords or tags. Then, retrieval systems could return a ranked list of items by computing the scores for all items w.r.t. a given user-query pair. However, the pre-targeting process could cost lots of computation resources. This significant challenge impulses us to think "Can we invent some algorithms to naturally represent the ternary interaction of $(query, user, item)$ in CR task?"

To answer this question, we next graphically review the relationship of $(user, query, item)$ and interpret how to model the ternary relationship from both users' and items' perspective.

3.2 Revisit Ternary Relationship

To formulate our ideal, we firstly represent the relationship of the $user \times query \times item$ triple as a graph, simply shown in Figure 1.(a), where entities $(user, query, item)$ are represented as vertices and red edges represent the observed relationship between a pair of entities. In order to intuitively describe the ternary relationship of $(user, query, item)$, we decompose Figure 1.(a) into two parts showed in Figure 1.(b) and Figure 1.(c), respectively. In Figure 1.(b), each positive observation $(user, query, item) \in \mathbf{X}$ can be represented as a route path *query-user-item*[5], where users serve as the intermedia vertices responsible for transforming the "resource" between items and queries, and the bipartite graph in the right intuitively represent the users' preferences over items. In terms of Figure 1.(c), the relationship of each positive observation is represented as a route path *user-item-query*, where differently items serve as the intermedia "resource" transformer between users and queries, and the left bipartite graph represent the rating information of user-item pairs. In the following part, we will give insight into the possible effects brought by the slight difference between Figure 1.(b) and (c).

[5]Here we use "path" to distinguish the representation of $(user, query, item)$ relationship from different perspectives.

It can be seen that Figure 1.(b) illustrates an user-central behaviour network, where the relationship of user-item pairs is represented as a collaborative network in the right part, analogously, the left part depicts the collaborative interactions of user-query pairs. In practical application, Figure 1.(b) could be extremely sparse due to the tinily available querying or rating information from users. In order to deal with such challenge caused by the lack of users' historical behaviours, the proposed approach for CR task should combine both collaborative networks in Figure 1.(b). By doing this, users' preferences over unconnected queries or items could be induced from their neighbours that have similar querying bias, or tend to rate a similar set of items.

Turning to Figure 1.(c), we can easily find that left part is shared with Figure 1.(b), but it represents a collaborative network from items' perspective. Analogous to Figure 1.(b), the right part of Figure 1.(c) represents the collaborative network of item-query pairs. Since content features of items might be the possible queries, the word "query" in Figure 1.(c) generally indicate both content features and the observed queries. With the problem of *asymmetric information* mentioned in Section 3.1, massive amount of items in Figure 1.(c) actually have too limited content to represent intrinsic feature. To better measure the ternary relationship of $(user, query, item)$, this challenge should be taken into account when we attempt to design an innovative algorithm for CR task. In CF field, many benefits of item-based methods have been fully discussed, among of which using the characteristics of neighbours with similar features to induce the target item's unknown features might be the possible way to help us to better evaluate the rank of sparse items w.r.t. a given user-query pair. However, the case in CR task is totally different from traditional CF problem, which usually involves binary relationship between user-item pairs, rather than the ternary interactions among $(user, query, item)$. As we mentioned in Section 3.1, pre-targeting items could cost a lot of computation in searching similar items. In this work, we prefer to combine both user-item and item-query collaborative networks to better capture the latent relationship of *user-item-query*. We believe that it could further tackle the sparsity problem caused by the lack of content features, meanwhile improve the effect of learning the ternary relevance of *user-item-query* path.

3.3 Our Method

As mentioned above, we consider items' collaborative information based on the assumption that regarding items as the media vertex could leverage the similar items [27] with rich descriptive terms to improve the performance of retrieval systems on estimating the ranks of sparse items. In our model, the score of *user-item-query* path in Figure 1.(c) can be given by:

$$g(q, a, u) = S_q A_a V_u^\top, \tag{2}$$

where $A_a \in \mathbb{R}^{n \times n}$ is the linear transformation matrix of item a. Finally, we propose to solve CR task by integrating the item-based collaborative information into LCR model, namely *Item-based latent collaborative retrieval* (ILCR). Consequently, Equation (1) can be modified as:

$$f(q, u, a) = S_q U_u T_a^\top + V_u T_a^\top + S_q A_a V_u^\top \tag{3}$$

where the first term intuitively captures the ternary interaction of $(query, user, item)$ by using encoder matrix U_u of each user to indirectly combine both user-query and user-item collaborative networks into one model. The querying and rating preferences of target users could be learned based on not only their own historical behaviours, but also their neighbours'. To some extent, it could effectively deal with the sparsity problem caused by the lack of users' historical information. Differently, the third term in Equation (3) focuses on modelling items' collaborative information to tackle the sparsity problem caused by the items' content features.

In terms of the scoring function, one might think that ILCR is slightly different from LCR, i.e. the performance of ILCR could be almost the same with LCR. But actually, the experimental results in Section 4.4 show that ILCR can better estimate the ranks of items w.r.t. a given user-query pair after considering the item-based collaborative information. The reason is that ILCR explicitly models item-based interactions in a sophisticated system that might be difficult to find for LCR method.

3.4 Pairwise Learning Approaches

The unknown parameters of the scoring function should be efficiently learned under the assumption that the top k retrieved items should include as many profitable selections as possible for a particular user w.r.t a given query. To achieve this, the CR task can be generally formulated as generating a ranked list of items for a given (u, q) pair by solving a pairwise ranking problem. A common approach is to regard each observation in a given training set $\mathbf{X}$, including m observations $(q_i, u_i, a_i)_{i=1,2,\ldots,m} \in \mathcal{Q} \times \mathcal{U} \times \mathcal{A}$, as a positive retrieval event, otherwise as a negative example. Subsequently, the selected learning algorithm should give the definition of *pairwise violation* cost, which would be produced if a negative item is assigned with a larger score or within a "margin" from the positive item.

In this section, we shall firstly detail the mechanism of WARP, originally used in the first piece of CR work [32], then introduce a generic learning algorithm proposed by Rendle et al. [22], namely BPR. Finally, we will show how to adapt BPR instead of WARP to correctly learn the f function of LCR and our model, ILCR.

3.4.1 Weighted Approximate-Ranking Pairwise (WARP)

The WARP Loss can be defined as follows:

$$err_{WARP} = \sum_{i=1}^{m} L(rank_{a_i}(\bar{f}(q_i, u_i))), \tag{4}$$

where $\bar{f}(q_i, u_i)$ is a vector, which contains predictions for all items for a fixed user-query pair. The a_i^{th} element of $\bar{f}(q_i, u_i)$, denoted as $\bar{f}_{a_i}(q_i, u_i)$, is the predicted relevance value of i^{th} training example (q_i, u_i, a_i). Correspondingly, $rank_{a_i}(\bar{f}(q_i, u_i))$ in Equation (4) is a margin-based rank of item a_i, defined as:

$$rank_{a_i}(\bar{f}(q_i, u_i)) = \sum_{b \neq a_i} \mathbb{I}[1 + \bar{f}_b(q_i, u_i) \geq \bar{f}_{a_i}(q_i, u_i)], \tag{5}$$

where $\mathbb{I}[\cdot]$ is the indicator function. In Equation (4), L is a weight function evaluating the loss of the current scoring function f:

$$L(k) = \sum_{i=1}^{k} \alpha_i, \ with \ \alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \ldots \geq 0. \tag{6}$$

According to the advices of [32], we choose $\alpha_i = 1/i$ as the weighting approach, which would assign large weights to top positions with rapidly decaying weight for lower positions. Intuitively, optimizing the WARP loss means to rank each positive item a_i in the training set to highest position. For example, given a random query-user pair denoted as (q_i, u_i), if the score of an uncollected item b is less than a margin of one from the score of a_i, this pair will produce a cost.

Subsequently, Equation (4) could be optimized by gradient descent based algorithms. However, in each updating step, it is expensive to compute the exact value of $rank_{a_i}$ for each observation when the number of items is very large. Thus, the exact rank of Equation (5) can be estimated by a random sampling process at each step [32]. That is to say, for a given observed sample (q_i, u_i, a_i), one uniformly draws at random items from $\mathcal{A}$ until finding a violated item b, which satisfies $1 + f(q_i, u_i, b) > f(q_i, u_i, a_i)$, and then the rank of a_i can be approximated as

$$rank_{a_i}(\bar{f}(q_i, u_i)) \approx \lfloor \frac{|\mathcal{A}| - 1}{K} \rfloor, \tag{7}$$

where $\lfloor \cdot \rfloor$ is the floor function and K is the number of steps needed to find a item b. Then Equation (4) can be modified as

$$\begin{aligned} \overline{err_{WARP}} &= \sum_{i=1}^{m} L_i \\ L_i &= L(rank_{a_i}(\bar{f}(q_i, u_i))) \cdot |1 - f(q_i, u_i, a_i) + f(q_i, u_i, b)|. \end{aligned} \tag{8}$$

Finally, Equation (8) could be optimized by stochastic gradient descent (SGD). We constrain the parameters using $||S_i|| \leq C, \ i \in \{1, \ldots, |\mathcal{Q}|\}$, $||V_i|| \leq C, \ i \in \{1, \ldots, |\mathcal{U}|\}$, $||T_i|| \leq C, i \in \{1, \ldots, |\mathcal{U}|\}$ and project the parameters back into the constraints at each SGD step.

3.4.2 Bayesian Personalized Ranking (BPR)

To clearly describe BPR algorithm [22], we leverage a notation D_X to denote the pairwise ranking constraints.

$$D_X = \{(u, q, a, b) : (u, q, a) \in \mathbf{X} \wedge (u, q, b) \notin \mathbf{X}\}$$

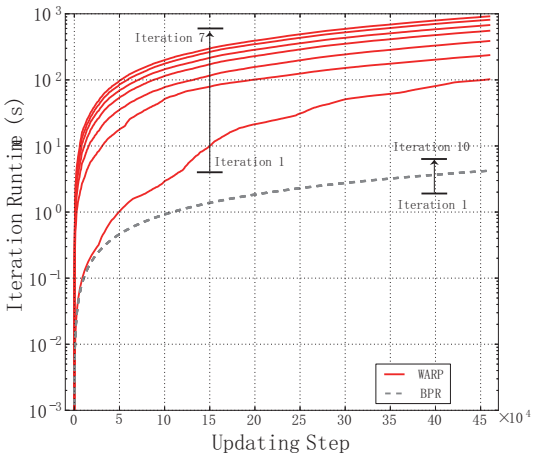

Figure 2: Accumulated runtime required by WARP and BPR for searching violated items in each training iteration. x axis denotes the number of updating steps in a training iteration. y axis denotes the accumulated runtime.

Next, we present a generic approach to solve the personalized ranking problem for CR tasks by maximizing the following posterior probability of the parameter space Θ.

$$p(\Theta|>_{u,q}) \propto p(>_{u,q}|\Theta)p(\Theta), \tag{9}$$

where $p(\Theta)$ denotes the prior probability of Θ, and notation $>_{u,q}= \{a >_{u,q} b : (u,q,a) \in \mathbf{X}, (u,q,b) \notin \mathbf{X}\}$ denotes the pairwise ranking structure for a given (u,q) pair. We assume that each element of $>_{u,q}$ is independently drawn from the same probability. Hence, the above likelihood function $p(>_{u,q}|\Theta)$ can be rewritten as:

$$p(>_{u,q}|\Theta) = \prod_{(u,q,a,b)\in D_X} p(a >_{u,q} b|\Theta), \tag{10}$$

where $p(a >_{u,q} b|\Theta)$ denotes the probability that a user really prefers item a to item b for a given query, defined as [22]:

$$\begin{aligned} p(a >_{u,q} b|\Theta) &= \sigma(\hat{x}_{u,q,a,b}(\Theta)) \\ \sigma(x) &= \frac{1}{1+e^{-x}}. \end{aligned} \tag{11}$$

Here we choose $\hat{x}_{u,q,a,b}(\Theta) = \hat{x}_{u,q,a} - \hat{x}_{u,q,b}$, in fact $\hat{x}_{u,q,a,b}(\Theta)$ could be a arbitrary real-valued function depending on the parameters Θ. For $p(\Theta)$, we define it as a normal distribution with zero mean and covariance matrix $\sum_\Theta = \lambda_\Theta I$, that is, $\Theta \backsim \mathcal{N}(\mathbf{0}, \sum_\Theta)$. Now we can infer the BPR-Opt by filling $p(\Theta)$ into the maximum posterior probability in Equation (9).

$$\begin{aligned} BPR\text{-Opt} &= \ln p(\Theta|>_{u,q}) \\ &= \ln \prod_{(u,q,a,b)\in D_X} p(a >_{u,q} b|\Theta)p(\Theta) \\ &= \ln \prod_{(u,q,a,b)\in D_X} \sigma(\hat{x}_{u,q,a,b}(\Theta))p(\Theta) \\ &= \sum_{(u,q,a,b)\in D_X} \ln \sigma(\hat{x}_{u,q,a,b}) + \ln p(\Theta) \\ &= \sum_{(u,q,a,b)\in D_X} \ln \sigma(\hat{x}_{u,q,a,b}) - \lambda_\Theta ||\Theta||^2, \end{aligned}$$

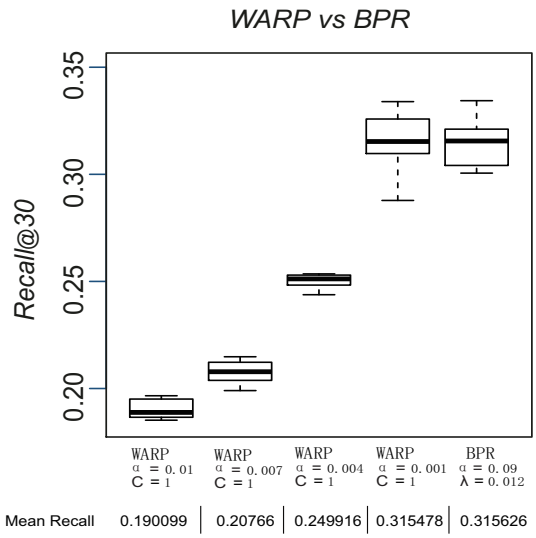

Figure 3: Comparisons of WARP and BPR on recall@30 (see Section 3.2). Four groups of (Learning rate α, constraint C) setups are randomly selected for WARP. One group of (Learning rate α, regularization λ) setup chosen for BPR mimics the best performance of WARP. The performance of each learning algorithm is validated on the same testing dataset (see Section 3.1).

where λ_Θ are model regularization parameters.

Typically, gradient ascent based algorithm is an apparent optimization strategy for maximizing the posterior probability in Equation (9). However, standard gradient ascent is not suitable for BPR-Opt because we have to compute $\hat{x}_{u,q,a,b}$ for all negative items b w.r.t. a given training sample $(u,q,a) \in \mathbf{X}$. If we have a large amount of items, it will be inefficient to update the parameters in each gradient ascent step. To effectively learn the parameters, a stochastic gradient-ascent (SGA) based algorithm, namely LearnBPR , was proposed by Rendle *et al.* [22] for optimizing BPR-Opt. Instead of comparing with all negative items, LearnBPR only draws at random a negative item b in each SGA step, which makes BPR superior to $WARP$ in terms of the computation complexity (see next section). The procedure of BPR is presented in **Algorithm 1**.

Algorithm 1 Optimizing models for BPR with LearnBPR

Input: $\mathbf{X}$: Training Set
Output: Learned $\hat{\Theta}$

```
initialize Θ
repeat
    randomly draw (u,q,a) from X
    randomly draw (u,q,b) from (u × q × A) \ X
    x̂_{u,q,a,b} ← x̂_{u,q,a} − x̂_{u,q,b}
    Θ ← Θ + α((1 − σ(x̂_{u,q,a,b})) ∂/∂Θ x̂_{u,q,a,b} − λ_Θ Θ)
until convergence.
```

3.5 Learning LCR and ILCR with BPR

In above sections, we intuitively review two optional learning algorithms for LCR model. In this section, we will present

the comparative results to study the efficiency of both learning algorithms.

In terms of WARP, the random sampling procedure dominates the runtime according to several recent works [31, 9]. Ref. [31] indicated that each SGD step requires less than $1 + min(\frac{|\mathcal{A}|-1}{rank_a(f(u,q))})$ sampling times to find a violated item b on average. Meanwhile, the computations of scores on items b could increase sharply or even worse for a massive item database, if the positive item a happens to rank at the top of the list. Interestingly, such occasional situation seems to happen frequently [9]. Comparatively, BPR only samples one time at each parameter updating step. In this paper, we conduct several experiments with the same hardware condition to demonstrate the efficiency of BPR for optimizing the LCR model. To avoid the fluctuations coming from parameters randomly initialization on the results, each experiment for Figure 3 runs ten independent times. Experimental results show that BPR needs less runtime to learn the parameters of LCR than WARP (see Figure 2), meanwhile, without decreasing the performance of LCR on the evaluation metric (see Figure 3). In Figure 2, the iteration runtime curves of BPR are difficult to be distinguished from each other, thus runtime curves of ten training iterations are crowded together.

In summary, we propose to employ BPR instead of WARP to optimize LCR model. The complete BPR learning procedure for LCR can be found in **Algorithm 2**, where the function of 5-12th steps is to update corresponding parameters based on the 5-6th steps in **Algorithm 1**.

Algorithm 2 Optimizing LCR with BPR

Input: $\mathbf{X}$: Training Set
Output: S, U, V, T

Initialize S, U, V, T from uniform probability $\mathcal{U}(x, y)$
repeat
  for (u,q,a) in $\mathbf{X}$ **do**
    randomly draw (u,q,b) from $(u \times q \times \mathcal{A}) \setminus \mathbf{X}$
    $\hat{x}_{u,q,a,b} \leftarrow f(u,q,a) - f(u,q,b)$
    Updating S_q, U_u, V_u, T_a, T_b
    $loss \leftarrow 1 - \sigma(\hat{x}_{u,q,a,b})$
    $S_q \leftarrow S_q + \alpha(loss \cdot U_u(T_a - T_b)^\top - \lambda S_q)$
    $V_u \leftarrow V_u + \alpha(loss \cdot (T_a - T_b) - \lambda V_u)$
    $T_a \leftarrow T_a + \alpha(loss \cdot (S_q U_u + V_u) - \lambda T_a)$
    $T_b \leftarrow T_b - \alpha(loss \cdot (S_q U_u + V_u) + \lambda T_b)$
    $U_u \leftarrow U_u + \alpha(loss \cdot S_q^\top (T_a - T_b) - \lambda U_u)$
  end for
until validation performance does not improve.

As Eequation (3) shows, the model parameters of ILCR include:

$$T \in \mathbb{R}^{|\mathcal{A}| \times n}, V \in \mathbb{R}^{|\mathcal{V}| \times n}, S \in \mathbb{R}^{|\mathcal{Q}| \times n}$$
$$A \in \mathbb{R}^{|\mathcal{A}| \times n \times n}, U \in \mathbb{R}^{|\mathcal{V}| \times n \times n}$$

For a given training dataset, the values of parameters in ILCR can be learned by following learning procedure of BPR. The complete learning procedure of ILCR is described in **Algorithm 3**, where the function of 5-14th steps is to update corresponding parameters based on the 5-6th steps in **Algorithm 1**.

Table 1: Basic statistics of $Lastfm$ and $Yelp$ dataset.

dataset	user	item	query	samples	sparsity
lastfm-50tags	1,529	8,669	50	574,521	99.91%
Yelp	16,826	14,902	587	806,261	99.99%

Algorithm 3 Optimizing ILCR with BPR

Input: $\mathbf{X}$: Training Set
Output: S, V, U, T, A

Initialize S, V, U, T, A from uniform probability $\mathcal{U}(x, y)$
repeat
  for (u, q, a) in $\mathbf{X}$ **do**
    randomly draw (u,q,b) from $(u \times q \times \mathcal{A}) \setminus \mathbf{X}$
    $\hat{x}_{u,q,a,b} \leftarrow f(u,q,a) - f(u,q,b)$
    Updating $S_q, V_u, T_a, T_b, U_u, A_a, A_b$
    $loss \leftarrow 1 - \sigma(\hat{x}_{u,q,a,b})$
    $S_q \leftarrow S_q + \alpha(loss \cdot (U_u(T_a - T_b)^\top + (A_a - A_b)V_u^\top) - \lambda S_q)$
    $V_u \leftarrow V_u + \alpha(loss \cdot (T_a - T_b + S_q(A_a - A_b)) - \lambda V_u)$
    $T_a \leftarrow T_a + \alpha(loss \cdot (S_q U_u + V_u) - \lambda T_a)$
    $T_b \leftarrow T_b - \alpha(loss \cdot (S_q U_u + V_u) + \lambda T_b)$
    $U_u \leftarrow U_u + \alpha(loss \cdot S_q^\top (T_a - T_b) - \lambda U_u)$
    $A_a \leftarrow A_a + \alpha(loss \cdot S_q^\top V_u - \lambda A_a)$
    $A_b \leftarrow A_b - \alpha(loss \cdot S_q^\top V_u + \lambda A_b)$
  end for
until validation performance does not improve.

4. EXPERIMENTAL RESULTS

In this section, we present experimental results to demonstrate the efficiency of the proposed collaborative retrieval model (ILCR) on two real-world datasets, one of which is obtained from the Last.fm music website, the other one is from *Yelp Dataset Challenge*[6] Round 3.

4.1 Dataset & Preprocessing

Last.fm dataset is released in the framework of HetRec 2011, called *hetrec*2011-*lastfm*-2*k* (*lastfm*-2*k*) [5], and contains heterogeneous information, which mainly covers users' collaborative tagging behaviours, listening preferences on artists, as well as social relationship.

In this paper, we focus on resources associated with collaborative retrieval tasks, that is, users' listening and tagging logs, which contain more than ten thousand unique tags utilized by 1892 users to reveal the characteristics of 17632 listened artists. While, these tags are used to represent artists with personal "notes" rather than widely accepted genres. To remove this noise, we keep only the top 50 most used tags, generally correlated to the genres of music: for example, the top 5 tags are "rock", "pop", "alternative", "electronic", and "indie".

To implement CR task, we need preprocess the dataset. It's noted that the query analysis applications will transform the provided queries as keywords or phrases which are the basic units to represent the items. Then retrieval system return a ranked list of items based on the parsed queries. Thereby, a common approach is to regard a *query*×*user*×*item* triple as a *keyword*×*user*×*item*, where the keywords are equivalent to the set of filtered user tags, or content features. We carry out data preprocessing with the following steps to

[6]http://www.yelp.com/dataset_challenge

Table 2: Parameters used for each method.

Method	$lastfm$-50$tags$	$Yelp$
ILCR	α=0.04, λ=0.01	α=0.08, λ=0.01
LCR	α=0.04, λ=0.01	α=0.1, λ=0.01
PITF	α=0.002, λ=0.01	α=0.02, λ=0.01

obtain an expectant dataset, denoted as $lastfm$-50$tags$, for the comparison experiments.

If an artist a has ever been listened and assigned with several tags by a user u, for example, *rock* and *indie*, then we allocate (*rock*, u, a) and (*indie*, u, a) to $lastfm$-50$tags$. If the artist is not assigned with any tags by the user u, the genres, assigned by other users to a, are distributed to the (u, a) pairs. If no user has ever assigned any genre to a, we choose to drop corresponding (u, a) pairs in default, since such user-artist pairs are not perfect training examples for the CR task. In the end, we infer the $lastfm$-50$tags$ (see Table 1) with 574,521 data points of the form (u, q, a) from the $last.fm$-2k dataset.

In addition to $last.fm$-2k dataset, this research is also performed on a recent academic dataset, published under the licence of *Yelp Dataset Challenge* Round 3, hereafter referred to as $Yelp$. This dataset contains 335,022 reviews and ratings given by 70746 users to 15470 businesses located in the Phoenix and AZ metropolitan area. Each business is characterized with rich content, for example, category like "Restaurants", "Shopping", "Health & Medical", which makes us accessible to have insight into users' preferences over different businesses. We pre-filter this data to contain users with at least 4 reviews, also corresponding businesses. After pre-filtering, we obtain a processed dataset including 806,261 well-formed points with format (user, category, business), given by 16826 users to 14902 businesses, which totally are pre-tagged with 587 business categories. For example, top-10 categories are "Restaurants", "Shopping", "Food", "Beauty & Spas", "Automotive", "Mexican", "Health & Medical", "Home Services", "Nightlife", "Fashion". Each sample (user, category, business) indicates that a particular user ever rated a business associated with the target category. Table 1 summarizes the basic information of the observed datasets.

It's noted that the sparsity of $lastfm$-50$tags$ is 99.91332%, approximately over 158 times denser than $Yelp$ dataset, which offers available stuff to distinguish the performance of our proposed model from other baseline algorithms in dealing with sparsity problem.

4.2 Experiment Settings

To demonstrate the efficiency of our proposed approach, we mainly compare ILCR with the state-of-the-art collaborative retrieval algorithm, that is, LCR model as well as algorithms for tensor environmental and traditional collaborative filtering. Besides LCR, the following baseline algorithms are used in this paper:

Pairwise interaction tensor factorization (PITF) [24]: As the state-of-the-art tensor algorithm for personalized tag recommendation, PITF is a personalized and context-aware algorithm aiming to return a top-k ranked list of tags when given a particular user-item pair. In PITF, the interaction between each entity pair is expressed as the dot-product of specific feature vectors. The ternary relationship $user \times item \times tag$ is modified by the following score function:

$$\hat{x}(u,t,i) = \hat{u}_u^T \hat{t}_t^U + \hat{u}_u^I \hat{i}_i^U + \hat{t}_t^I \hat{i}_i^T, \tag{12}$$

where the first term $\hat{u}_u^T \hat{t}_t^U$ denotes the relevance value of the given user u and tag t, middle term $\hat{u}_u^I \hat{i}_i^U$ denotes the relevance value of the given user u and item i, and the last term $\hat{t}_t^I \hat{i}_i^T$ denotes the interaction between the given item i and tag t. According to [24], the user-item interaction term vanishes when LearnBPR is employed to optimize top-k ranking task. Thus, the final score function for tag recommendation task is:

$$\hat{x}(u,t,i) = \hat{u}_u \hat{t}_t^U + \hat{i}_i \hat{t}_t^I. \tag{13}$$

Analogously, collaborative retrieval task involves with returning a top-k ranked list of items w.r.t. a given user-query pair, which inspires us to adapt it to CR tasks by lightly modifying Equation (12). Likewise, the user-query interaction term vanishes, then the modified score function is denoted as:

$$\hat{x}(u,q,i) = \hat{u}_u \hat{i}_i^U + \hat{q}_q \hat{i}_i^Q. \tag{14}$$

where the first term $\hat{u}_u \hat{i}_i^U$ denotes the interaction between user u and item i, and the last term denotes the interaction between query q and item i. The values of parameters can be leaned by employing BPR algorithm [24].

Non-negative Matrix Factorization (NMF) [18]: This approach is not directly applied to model ternary relationship. In this paper, we perform NMF on the item-query matrix to compute a top-k ranked preference items for given q and u. The NMF implementation we used in this paper is from http://www.csie.ntu.edu.tw/∼cjlin/nmf/.

In this paper, we conduct numerous experiments to evaluate the performances of each algorithm on the real Last.fm and Yelp datasets. In terms of Last.fm dataset, we randomly draw 80% of samples in $last.fm$-50$tags$ for training, 10% for validation, and the rest for testing. Based on this dataset, we study the efficiency of WARP and BPR on optimizing LCR model (see **Section 3.4**). The performance is evaluated on the testing dataset. In addition, we compare ILCR with LCR to validate the capacity of each CR model on retrieving items with different size of training samples (40%, 60%, 80%, 100% of the total training data). Analogously, we randomly draw 60% of samples in $Yelp$ for training, 20% for validation, and the rest for testing. The initial values of feature matrices of both models are randomly drawn from a uniform distribution $\mathcal{U}(-0.02, 0.02)$. The hyperparameters α and λ are chosen for each algorithm using the validation set respectively. Different experiment settings are used based on the dataset and algorithms (see Table 2).

4.3 Evaluation

The performance of each algorithm is measured by recall@k, a widely used metric to evaluate the recommendation accuracy in top-k. For a given testing example (q,u,a), we first

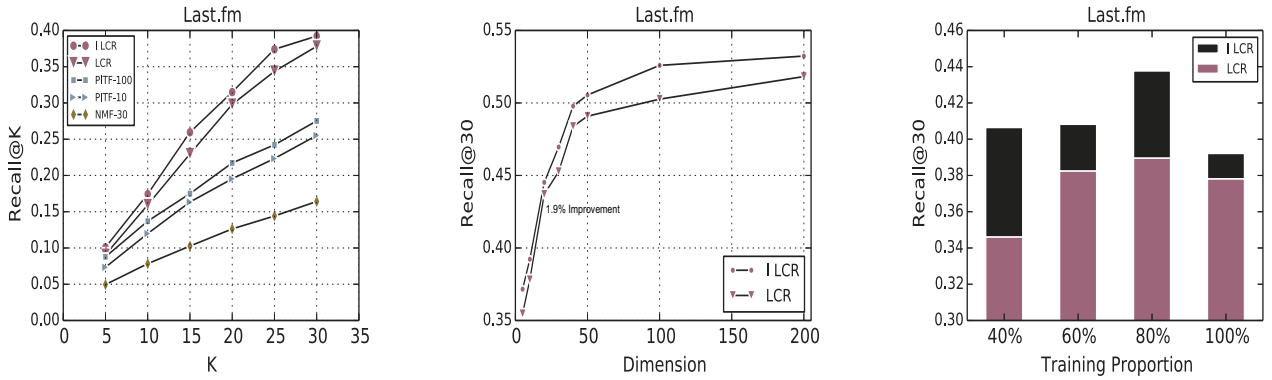

(a) (b) (c)

Figure 4: Experimental results on $Last.fm$ dataset. (a) Different algorithms' *Recall* vs. different k. The integer numbers of PITF and NMF represent the dimension n, and n is set as 10 for ILCR and LCR algorithms. (b) Recall@30 vs. the dimension parameter n. (c) Recall@30 vs. the dimension parameter n.

compute $f(q, u, i)$ for each item $i \in \mathcal{A}$, and then sort them in descending order of score. Then recall@k equals to 1 if artist a appears in the top k list, and equals to 0 otherwise. We report mean recall@k over the whole test dataset.

4.4 Results

The experimental results on the Last.fm dataset are depicted in Figures 4. From Figure 4.(a), we can see that the correctly predicted items of both ILCR and LCR increase as the growth of different values of k with fixed dimension n=10. Comparing with other baseline algorithms, ILCR performs best in retrieving items with given queries by users. Noted that, in comparison with LCR in Figure 4.(a), ILCR increases at least 3.7% when k=30, whilst improves over 12% in the best case when k equals to 15. In terms of PITF, recall ratio improves along with the increase of both parameters k and dimension n. However, PITF still obtains lower recall value than LCR and ILCR on the Last.fm dataset. The performance of NMF suggests that only taking into account binary relationship is insufficient for collaborative retrieval task. Figure 4.(b) shows that the performance of both ILCR and LCR improve along with the increase of dimension n, while ILCR outperforms LCR with over 1.9% improvement when $n = 20$. In addition to evaluating the efficiency of involved algorithms with different chosen values of k and n, the performance of ILCR and LCR on various size of training samples is also shown in Figure 4.(c). The sparsity of training dataset declines as the growth of the size of training samples. Figure 4.(c) reveals that ILCR could effectively represent the latent relationship of the ($user$,$query$,$item$) triples, even for different sparse datasets.

Since we expect the integration of item-based information to be particularly useful when the dataset includes massive amount of sparse items, we also show experimental results on a much sparser dataset, $Yelp$. Figure 5.(a) shows that the recall value of both ILCR and LCR improve with the increase of parameter k. From Figure 5.(a) and (b), we can see that the performance of LCR on the $Yelp$ dataset significantly decreases with different settings, in contrast with the well performance on the $Last.fm$ dataset. It's noted that we carefully and repeatedly choose the best setting of key hyperparameters α and λ on the validation dataset. However, the difference of recall ratio between LCR and ILCR is still evident, which suggests that ILCR can be useful for evaluating the ranks of items, especially those sparse items. Figure 5.(a) shows that ILCR outperforms the selected baseline algorithms at different values of k with fixed dimension parameter n=10. The experimental results demonstrate that taking item's collaborative information into account can effectively predict users' preferences on items w.r.t. a given query under a situation where there is imperfect knowledge of the characteristics of the majority proportion of items.

5. RELATED WORK

Recommendation and retrieval techniques have become essential components of massive applications, like E-commerce, search engine, music-based social network etc. In information retrieval, one is expected to rank items or documents by using the content features of items depending on a given query by an active user. In that case, many algorithms are proposed, like Latent Semantic Indexing [6], LDA [4] topic model for low-dimensional representation of the word. More recently, factorized models like Polynomial Semantic Indexing (PSI) [2] are employed to implement the task of document retrieval. Methods like PSI for retrieval task optimizes the AUC ranking loss, which does not optimise the top k ranked list like ours. In recommendation field, many works based on factorized models are proposed to produce predictions to active users. In particular, Matrix Factorization-based methods [17, 15, 26] are very close to our method since each entity (e.g. user, item, query) is represented as a low-dimensional feature vector. Main difference between ILCR and these general matrix factorization methods is that each recommendation in CR task is seeded with a query, while most of them are invented to deal with binary user-item relationship. In fact, the most related works to our model are those models which can be adaptive to deal with multi-types of relationship as ternary interactions are considered in CR task. There are several available choices, such as classical Tucker decomposition [30], PARAFAC. Many context-aware collaborative filtering techniques are also proposed for such multi-relationship learning task, in particular contextual information related to users, like tags [24], web pages [21],

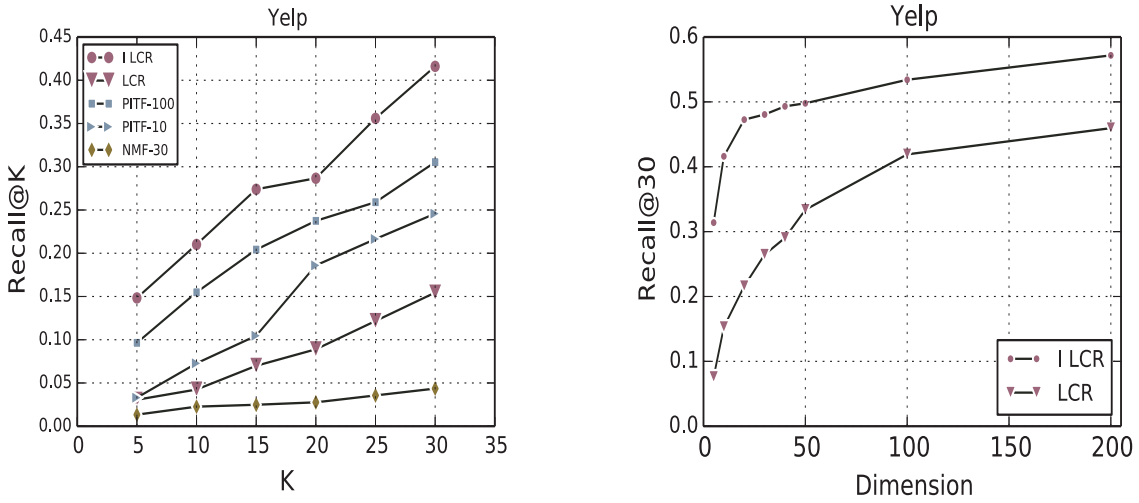

(a) (b)

Figure 5: Experimental results on $Yelp$ dataset. (a) Different algorithms' *Recall* vs. different k. The integer numbers of PITF and NMF represent the dimension n, and n is set as 10 for ILCR and LCR algorithms. (b) Recall@30 vs. the dimension parameter n.

demographics [14], mobile information [35], temporal effects [33]. In addition, recent works on heterogeneous information network [28, 34] can also be adaptive to give vision into the sophisticated interactions hidden in the complicated dataset.

6. CONCLUSIONS AND FUTURE WORKS

We aim to design an effective collaborative retrieval algorithm to objectively predict the ranks of those items with lack of descriptive terms revealing the basic characteristics of them. To achieve this goal, we focus on the profits of leverage the collaborative information of items to better evaluate the ranks of those sparse items, and propose to express the latent sophisticated relationships for CR task from not only users' perspective, but also items' perspective. Then, we propose a superior latent collaborative retrieval model, ILCR, by integrating the possible item-based information into LCR model. Finally, we study the primary principles of LCR, of which we find that the runtime of LCR is dominated by the chosen learning approach, namely WARP. Therefore, we propose to employ a generic approach, namely BPR, instead of WARP to optimize the LCR model on the basis of further experimental analysis covering the topics of training efficiency and accuracy.

The proposed model can be easily generalized to deal with many other tasks involving to model ternary interaction among entities, not limited to CR task. Most typical example might be the tagging recommendation system, which aims to recommend tags to users w.r.t. items. However, we just explore possible interactions among $(user, query, item)$ triple. In practice, relationships of a pair of entities always involve with massive ingredients, which could be termed as a currently-prevalent word, heterogeneous relationships. Due to this significant characteristics, we will attempt to make our proposed model adaptive to learn the multi-relation by exploring diverse aspects of a sophisticated system in future work.

7. ACKNOWLEDGEMENTS

This work was partially supported by the National Natural Science Foundation of China (No. 11305043), and the Zhejiang Provincial Natural Science Foundation of China (No. LY14A050001), the EU FP7 Grant 611272 (project GROWTHCOM) and Zhejiang Provincial Qianjiang Talents Project (Grant No. QJC1302001).

8. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In *Recommender systems handbook*, pages 217–253. Springer, 2011.

[2] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasa, Y. Qi, C. Cortes, and M. Mohri. Polynomial semantic indexing. In *NIPS'09*, pages 64–72, 2009.

[3] L. Baltrunas, B. Ludwig, and F. Ricci. Matrix factorization techniques for context aware recommendation. In *Proceedings of the ACM conference on Recommender systems*, pages 301–304. ACM, 2011.

[4] D. M. Blei, A. Y. Ng, M. I. Jordan, and J. Lafferty. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:2003, 2003.

[5] I. Cantador, P. Brusilovsky, and T. Kuflik. Second workshop on information heterogeneity and fusion in recommender systems (hetrec2011). In *RecSys*, pages 387–388, 2011.

[6] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41:391–407, 1990.

[7] M. Deshpande and G. Karypis. Item-based top- n recommendation algorithms. *ACM Transactions on Information Systems*, 22:143–177, 2004.

[8] N. Garg and I. Weber. Personalized, interactive tag

recommendation for flickr. In *Proceedings of the ACM conference on Recommender systems*, pages 67–74. ACM, 2008.

[9] K.-J. Hsiao, A. Kulesza, and A. Hero. Social collaborative retrieval. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 293–302, 2014.

[10] M. Jamali and M. Ester. Trustwalker: a random walk model for combining trust-based and item-based recommendation. In *KDD'09*, pages 397–406, 2009.

[11] R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in folksonomies. In *Knowledge Discovery in Databases: PKDD 2007*, pages 506–514. Springer, 2007.

[12] B. Kanagal, A. Ahmed, S. Pandey, V. Josifovski, L. Garcia-Pueyo, and J. Yuan. Focused matrix factorization for audience selection in display advertising. In *Proceedings of the 2013 IEEE international conference on data engineering (ICDE)*, pages 386–397, 2013.

[13] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 79–86. ACM, 2010.

[14] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *RecSys'10*, pages 79–86, 2010.

[15] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *KDD'08*, pages 426–434, 2008.

[16] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.

[17] Y. Koren, R. M. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *IEEE Computer*, pages 30–37, 2009.

[18] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[19] G. Linden, B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, pages 76–80, 2003.

[20] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 287–296. ACM, 2011.

[21] A. K. Menon, K. P. Chitrapura, S. Garg, D. Agarwal, and N. Kota. Response prediction using collaborative filtering with hierarchies and side-information. In *KDD'11*, pages 141–149, 2011.

[22] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on Uncertainty in Artificial Intelligence*, pages 452–461. AUAI Press, 2009.

[23] S. Rendle, L. B. Marinho, A. Nanopoulos, and L. Schmidt-thieme. Learning optimal ranking with tensor factorization for tag recommendation. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 727–736, 2009.

[24] S. Rendle and L. Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 81–90. ACM, 2010.

[25] S. Riedel, L. Yao, and A. McCallum. Relation extraction with matrix factorization and universal schemas. In *Joint Human Language Technology Conference/Annual Meeting of the North Computational Linguistics (HLT-NAACL)*, June 2013.

[26] R. Salakhutdinov and A. Mnih. Probabilistic Matrix Factorization. In *Neural Information Processing Systems*, 2007.

[27] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295. ACM, 2001.

[28] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. In *PVLDB*, pages 992–1003, 2011.

[29] P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos. Tag recommendations based on tensor dimensionality reduction. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 43–50, 2008.

[30] L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31:279–311, 1966.

[31] J. Weston, S. Bengio, and N. Usunier. Large scale image annotation: learning to rank with joint word-image embeddings. *Machine learning*, 81(1):21–35, 2010.

[32] J. Weston, C. Wang, R. Weiss, and A. Berenzweig. Latent collaborative retrieval. In *Proceedings of the 29th International Conference on Machine Learning*, pages 9–16, 2012.

[33] L. Xiong, X. Chen, T.-K. Huang, J. Schneider, and J. G. Carbonell. Temporal Collaborative Filtering with Bayesian Probabilistic Tensor Factorization. In *SIAM International Conference on Data Mining*, pages 211–222, 2010.

[34] X. Yu, X. Ren, Y. Sun, Q. Gu, B. Sturt, U. Khandelwal, B. Norick, and J. Han. Personalized entity recommendation: a heterogeneous information network approach. In *WSDM'14*, pages 283–292, 2014.

[35] V. W. Zheng, B. Cao, Y. Zheng, X. Xie, and Q. Yang. Collaborative Filtering Meets Mobile Recommendation: A User-Centered Approach. In *National Conference on Artificial Intelligence - AAAI*, 2010.